arXiv:1702.05141v2 [math.CO] 3 Mar 2017

L-INFINITY OPTIMIZATION TO BERGMAN FANS OF MATROIDS WITH AN APPLICATION TO PHYLOGENETICS

DANIEL IRVING BERNSTEIN

Abstract. Given a dissimilarity map δ on finite set X, we show that the set of ultrametrics (equidistant tree metrics) which are l^∞-nearest to δ is a tropical polytope. We give an algorithm for computing a superset of its tropical vertices. It was shown by Ardila and Klivans that the set of all ultrametrics on a finite set of size n is the Bergman fan associated to the matroid underlying the complete graph on n vertices. Therefore, we derive our results in the more general context of Bergman fans of matroids. This more general context allows our algorithm to be applied to dissimilarity maps where only a subset of the entries are known.

1. Introduction

A fundamental problem in phylogenetics is to infer the evolutionary history among a collection of species from data. One approach is to use *distance-based methods*. The data required for such an approach is some measure of distance between each pair of species. If these distances are computed using some property that is expected to change in proportion to time elapsed, then one often assumes that the pairwise distances approximate an *ultrametric*. Finding a best-fit ultrametric to an arbitrary dissimilarity map is therefore an important computational problem.

In tropical geometry, one redefines arithmetic over the real numbers so that the sum of two numbers is their maximum and the product is their sum (in the usual sense). There are strong connections between phylogenetics and tropical geometry [2, 3, 16, 13, 12] so the l^∞-metric is a natural choice to measure best fit for phylogenetic reconstruction.

An algorithm of Chepoi and Fichet computes an l^∞-nearest ultrametric to a given dissimilarity map in polynomial time [5] but this is generally not the only l^∞-nearest ultrametric. In fact, the set of l^∞-nearest ultrametrics may be high-dimensional and represent many different tree topologies [4]. So from the perspective of the l^∞-metric, there may be many different evolutionary histories that all explain a given dataset equally well. The main motivation of this paper is to understand the structure of the set of l^∞-nearest ultrametrics and our first result towards this is the following.

Proposition 3.2. *Let δ be a dissimilarity map on a finite set X. The set of ultrametrics that are nearest to δ in the l^∞-norm is a tropical polytope.*

A tropical polytope is the tropical convex hull of a finite set of points. All ultrametrics that are l^∞-nearest to a given dissimilarity map have the same underlying tree structure (called the *topology*) if and only if all tropical vertices of the tropical polytope they form do as well. We give an algorithm for computing the tropical vertices of the tropical polytope

described in Proposition 3.2.

We derive our results in a more general context. Ardila and Klivans showed that the set of ultrametrics on n species is the Bergman fan associated to the matroid underlying the complete graph on n vertices [3]. Therefore we can view the problem of finding the set of l^∞-nearest ultrametrics as a special case of the problem of finding the set of l^∞-nearst points in the Bergman fan of a matroid. We have the following generalization of Proposition 3.2.

Proposition 4.5. *If $\mathcal{M}$ is a matroid on ground set E and $x \in \mathbb{R}^E$, then the subset of the Bergman fan of $\mathcal{M}$ consisting of elements l^∞-nearest to x is a tropical polytope.*

The main result of this paper is an algorithm for computing the tropical vertices of the tropical polytope described in Proposition 4.5.

Generalizing to Bergman fans of arbitrary matroids has a potential application in phylogenetics. Namely, if one wishes to reconstruct a phylogeny from partial distance data where observed distances correspond to the edges of some graph G, then one can begin by optimizing to the Bergman fan of G's matroid which will give a partial ultrametric (see Proposition 4.32). This reconstruction problem is a special case of the *sandwich to ultrametric problem* studied by Farach, Kannan, and Warnow in [7].

This paper is organized as follows. Section 2 gives the necessary background on tropical convexity. Section 3 gives our algorithm for computing a finite generating set of the tropical polytope consisting of the ultrametrics that are l^∞-nearest to a given dissimilarity map. Section 4 generalizes the concept of an ultrametric's tree topology to elements of Bergman fans of arbitrary matroids and uses this to generalize the algorithm given in Section 3. Section 5 applies our algorithm to a biological dataset.

Acknowledgments

The author is grateful to Colby Long and Seth Sullivant for many helpful conversations and for feedback on early drafts. The author was partially supported by the US National Science Foundation (DMS 0954865) and the David and Lucille Packard Foundation.

2. Preliminaries on Tropical Convexity

This section reviews the necessary concepts from tropical convexity. There are at least two different sets of basic definitions related to tropical convexity. One is used in [6], and the other in [1]. We adhere to the conventions of the latter, as their definition of tropical polytope is more natural in our context.

The *tropical semiring*, also known as the max-plus algebra, is the set $\mathbb{R} \cup \{-\infty\}$ together with the operations $a \oplus b := \max\{a, b\}$ and $a \odot b := a + b$. We denote this semiring by $\mathbb{R}_{\max}$. The additive identity of $\mathbb{R}_{\max}$ is $-\infty$ and the multiplicative identity is 0. The set $\mathbb{R}^n_{\max}$ is an $\mathbb{R}_{\max}$-semimodule where for $x, y \in \mathbb{R}^n_{\max}$ and $\alpha \in \mathbb{R}_{\max}$, $(x \oplus y)_i := x_i \oplus y_i$ and $(\alpha \odot x)_i := \alpha + x_i$. If $A \in \mathbb{R}^{m \times n}_{\max}$ is a matrix and $x \in \mathbb{R}^n_{\max}$, then the product $A \odot x$ is the usual matrix product, but with multiplication and addition interpreted tropically. That

is, if A has columns $a_1, \ldots, a_n$, then

$$A \odot x := \bigoplus_{j=1}^{n} x_j \odot a_j.$$

Several notions from ordinary convexity theory have tropical analogs. We say that $P \subseteq \mathbb{R}_{\max}$ is a *tropical cone* if whenever $x, y \in P$ and $\lambda, \mu \in \mathbb{R}_{\max}$, $\lambda \odot x \oplus \mu \odot y \in P$. If this only holds with the restriction that $\lambda \oplus \mu = 0$, then we say that P is *tropically convex*.

A *tropical polyhedron* is a set of the form

$$\{x \in \mathbb{R}^n_{\max} : A \odot x \oplus b \geq C \odot x \oplus d\}$$

where $A, C \in \mathbb{R}^{m \times n}_{\max}$ and $b, d \in \mathbb{R}^m_{\max}$. We denote this set $P(A, b, C, d)$. Note that $P(A, b, C, d)$ is always tropically convex. When $b = d = (-\infty, \ldots, -\infty)^T$ then $P(A, b, C, d)$ is a tropical cone and we call it a *tropical polyhedral cone*. Bounded tropical polyhedra are called *tropical polytopes*. Given $V \subseteq \mathbb{R}_{\max}$, $\operatorname{tconv}(V)$ is the *tropical convex hull* of V. That is,

$$\operatorname{tconv}(V) := \{\lambda \odot x + \mu \odot y : x, y \in V, \lambda \oplus \mu = 0\}.$$

We define the *tropical conic hull* $\operatorname{tcone}(V)$ similarly. Gaubert and Katz showed in [9] that any tropical polytope (cone) P can be expressed as the tropical convex (conic) hull of a finite set V. Conversely, Joswig showed in [11] that if $V \subseteq \mathbb{R}^n$ is a finite set and $P = \operatorname{tcone}(V)$, then P is a tropical polyhedral cone. The analogous result for $P = \operatorname{tconv}(V)$ follows from results in [9]. We call the minimal such V the *tropical vertices (extreme rays) of* P.

3. Results for phylogenetics: l-infinity nearest ultrametrics

This section gives our algorithm for finding a generating set of the tropical polytope consisting of the ultrametrics that are l^∞-nearest to a given dissimilarity map. This section presents the results of Section 4 in the context of our main motivation. We begin by reviewing the necessary background about ultrametrics, which are a special type of tree metric. We then give our algorithm which we illustrate on an example.

Let $X = \{x_1, \ldots, x_n\}$ be a finite set. A *dissimilarity map* on X is a function $\delta : X \times X \to \mathbb{R}$ such that $\delta(x, x) = 0$ and $\delta(x, y) = \delta(y, x)$ for all $x, y \in X$. We can express a dissimilarity map d as a matrix D where $D_{ij} = \delta(x_i, x_j)$. Note that D is symmetric with zeros along the diagonal. Such a matrix will be called a *dissimilarity matrix*. A *rooted X-tree* T is a tree whose leaf set is X and that has exactly one vertex of degree 2, called the *root of* T denoted $\operatorname{root}(T)$. A *descendant* of a vertex v in a rooted tree T is a node $u \neq v$ such that the unique path from u to $\operatorname{root}(T)$ contains v. Note that all non-root vertices are descendants of $\operatorname{root}(T)$. The set of descendants of a vertex v in a rooted tree T is denoted $\operatorname{Des}_T(v)$. We let T° denote the collection of interior vertices of T. Let T be a rooted X-tree and let $w : T^\circ \to \mathbb{R} \setminus \{0\}$ be a weighting of the internal nodes of T such that $w(u) \leq w(v)$ whenever $u \in \operatorname{Des}_T(v)$. The pair (T, w) gives rise to a dissimilarity map $\delta_{T,w}$ on X such that $\delta_{T,w}(x_i, x_j) = w(v)$ where $v \in T^\circ$ is the vertex nearest to $\operatorname{root}(T)$ in the unique path from x_i to x_j. Given a dissimilarity map δ on X,

if we can express δ as $\delta_{T,w}$ for some X-tree T with internal node weights w, then we say that δ is an *ultrametric*. If we require that $w(u) < w(v)$ whenever $u \in \mathrm{Des}_T(v)$, then the rooted X-tree T is unique and we call it the *topology* of δ. Figure 1 shows an ultrametric δ along with an interior-vertex-weighted tree displaying δ.

A *polytomy* of a rooted tree is either a non-root internal node of degree at least four, or the root node if it has degree at least three. We say that a rooted tree is *binary* if it does not have any polytomy. A *resolution* of a tree T is a binary tree T' such that T can be obtained from T' via a (possibly empty) series of edge contractions. A resolution of an ultrametric δ is a rooted tree representation (T, w) of δ such that T is binary. Note that if the topology underlying δ is not binary, then there will be multiple resolutions of δ. Figure 2 illustrates these concepts by representing a single ultrametric in three ways - on its topology and on two different resolutions.

$$\begin{array}{c} \\ A \\ B \\ C \\ D \end{array} \begin{array}{c} \begin{array}{cccc} A & B & C & D \end{array} \\ \begin{pmatrix} 0 & 5 & 7 & 9 \\ 5 & 0 & 7 & 9 \\ 7 & 7 & 0 & 9 \\ 9 & 9 & 9 & 0 \end{pmatrix} \end{array}$$

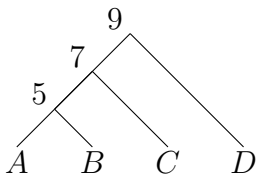

FIGURE 1. An ultrametric on $\{A, B, C, D\}$ and its representation on a rooted tree

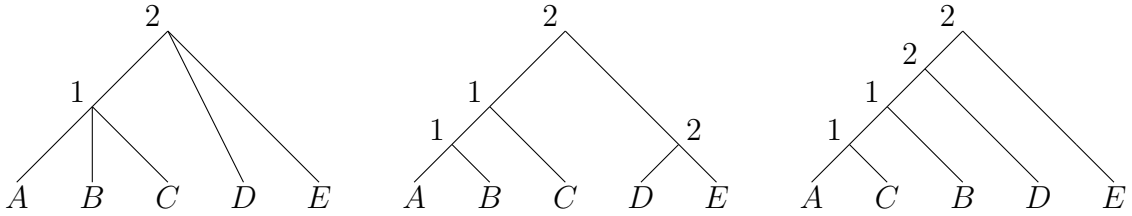

FIGURE 2. An ultrametric whose topology has two polytomies represented on its topology and two different resolutions

Given two dissimilarity maps δ_1, δ_2 on X with associated distance matrices D_1, D_2, we define the l^∞ *distance between* δ_1 *and* δ_2, denoted $\|\delta_1 - \delta_2\|_\infty$, to be the greatest absolute value among entries in $D_1 - D_2$. An important question that comes up in phylogenetics is then: given a dissimilarity map δ, which ultrametrics are nearest to δ in the l^∞ metric? Chepoi and Fichet [5] give an algorithm for producing a single ultrametric l^∞-nearest to a given dissimilarity map which we now describe.

Theorem 3.1 ([5])**.** *Let δ be a dissimilarity map on a finite set X. Then the following algorithm produces an ultrametric on X that is nearest to δ in the l^∞ norm.*

(1) *Draw the complete graph on vertex set X*
(2) *Label the edge between x and y by $\delta(x, y)$*

(3) *For each* $x, y \in X$ *define*

$$\delta_u(x,y) = \min_{\text{paths } P \text{ from } x \text{ to } y} \left(\max_{\text{edges } (i,j) \text{ of } P} \delta(i,j) \right)$$

(4) *Let* $d = \|\delta_u - \delta\|_\infty$ *and let* $\mathbf{1}$ *be the ultrametric such that* $\mathbf{1}(x,y) = 1$ *for all* $x, y \in X$
(5) *Then* $\delta_u(x,y) + \frac{d}{2}\mathbf{1}$ *is an ultrametric that is* l^∞*-nearest to* δ.

Although the algorithm given by Theorem 3.1 produces only one ultrametric, there can be many ultrametrics that are l^∞-nearest to a given dissimilarity map and they can have different topologies. For example, the dissimilarity map given in Figure 3 is l^∞-nearest to many different ultrametrics. Two with differing topologies are shown. We call the ultrametric given by Theorem 3.1 the *maximal closest ultrametric to* δ and denote it symbolically as δ_m. That δ_m is indeed maximal among ultrametric nearest to δ is shown in [5], and also follows from Lemma 4.15(4).

$$\begin{array}{c} \\ A \\ B \\ C \\ D \end{array} \begin{array}{c} \begin{array}{cccc} A & B & C & D \end{array} \\ \begin{pmatrix} 0 & 2 & 4 & 6 \\ 2 & 0 & 8 & 10 \\ 4 & 8 & 0 & 12 \\ 6 & 10 & 12 & 0 \end{pmatrix} \end{array}$$

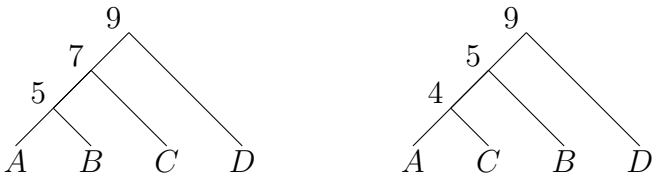

FIGURE 3. A dissimilarity map and two l^∞-nearest ultrametrics with different topologies

Proposition 3.2. *Let* δ *be a dissimilarity map on a finite set* X*. The set of ultrametrics that are nearest to* δ *in the* l^∞*-norm is a tropical polytope.*

We will prove Proposition 3.2 in a more general setting later. We now give an algorithm for computing the tropical vertices of the set of l^∞-nearest ultrametrics to a given dissimilarity map. Correctness and termination will be proven in the next section in a more general context.

The basic idea behind our algorithm is as follows. First, we compute δ_m and represent it on its topology with the internal nodes weighted. At each iteration, we choose a previously identified l^∞-nearest ultrametric and some internal node of its topology. We then decrease the weight on that node until either a polytomy is induced, or decreasing it any further would increase the distance between the resulting ultrametric and δ. The resulting ultrametric is then kept and included as part of the output. This is described precisely in Algorithm 1 and illustrated in Example 3.3. It is worth noting that all ultrametrics l^∞-nearest to a given dissimilarity map δ have the same topology if and only if the ultrametrics in the output of Algorithm 1 all have the same topology. This follows from Corollary 4.29.

Example 3.3. Let δ be the ultrametric given in Figure 3 on the left. We trace through Algorithm 1 when given input δ. We will make reference to ultrametrics $u_1, \ldots, u_5$ which are shown in Figure 4. The set S_0 contains only u_1. The ultrametric u_1 is binary so it

Algorithm 1 Compute a generating set of the set of ultrametrics l^∞-nearest to δ

Input: A dissimilarity map $\delta : X \times X \to \mathbb{R}$
  Compute δ_m using Theorem 3.1
  Set $S_0 := \{\delta_m\}$, $S_i := \emptyset$ for $i = 1, 2, \ldots$
  for $i := 1, 2, \ldots$ **do**
    if $S_{i-1} = \emptyset$ **then**
      break
    end if
    for $u \in S_{i-1}$ **do**
      for each resolution (T, w) of u **do**
        for each internal vertex v of T such that $w(v) > w(u)$ for all $u \in \mathrm{Des}_T(v)$ **do**
          set $u := \arg\max_{u \in \mathrm{Des}_T(v)} w(u)$
          if setting $w(v) = w(u)$ produces an ultrametric l^∞-nearest to δ **then**
            reset $w(v) := w(u)$
          else
            reset $w(v)$ smallest possible so that $\delta_{T,w}$ is still l^∞-nearest to δ
          end if
          if $u \neq \delta_{T,w}$ **then**
            add $\delta_{T,w}$ to S_i
          end if
        end for
      end for
    end for
  end for
  return $\bigcup_{i=0}^{\infty} S_i$

has only one resolution. Let v be the node of u_1's topology with weight 5. The smallest number we can replace this with while maintaining l^∞-distance from δ is -1, so we change that weight. The resulting ultrametric is u_3 which gets added to S_1. Alternatively, letting v be the vertex with weight 7, we can change the weight of v to 5 giving us the ultrametric u_2 which also gets added to S_1.

As S_1 is nonempty, the algorithm continues. Choosing $u = u_2$, we have two possible resolutions. Resolving to give us the topology of u_4 allows us to give the new node a weight of 1 without increasing l^∞ distance, and nothing smaller. This gives us u_4 which we then include in S_2. If we resolve to get the topology of u_5, the smallest weight we can give the new node is -1, and relabeling as such gives us u_5 which we then add to S_2. Choosing $u = u_3$ gives us only one resolution as u_3 has a binary topology. The only vertex that can be reset as described in the algorithm while giving a new tree is the vertex with weight 7. We can make the label on this vertex as small as 5 without increasing l^∞ distance from δ.

Now S_2 is nonempty so the algorithm continues. However, there is no way to decrease vertex weights as the algorithm describes without increasing l^∞ distance to δ and so S_3 is empty and the algorithm terminates. Correctness of the algorithm implies that the

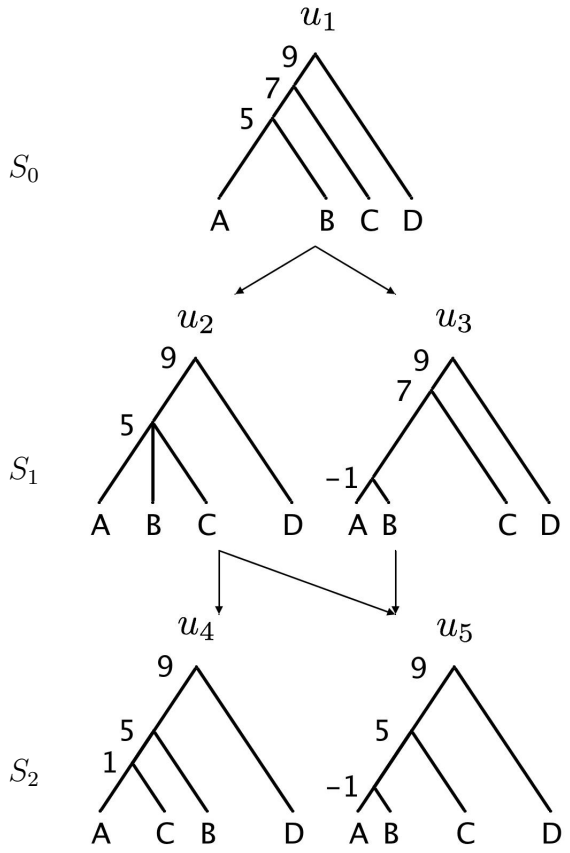

FIGURE 4. Ultrametrics produced by Algorithm 1 when input is dissimilarity map given on the left side of Figure 3

set of ultrametrics l^∞-nearest to δ is $\text{tconv}\{u_1, \ldots, u_5\}$. This generating set is redundant because $u_2 = u_4 \oplus u_5$ and $u_1 = u_3 \oplus u_4$. In fact, the tropical vertices of this set are u_3, u_4, and u_5. Note that this tropical polytope is contained in the three-dimensional affine subspace $\{\tilde{\delta} \in \mathbb{R}^{\binom{[4]}{2}} : \tilde{\delta}(1,4) = \tilde{\delta}(2,4) = \tilde{\delta}(3,4) = 9\} \subset \mathbb{R}^{\binom{[4]}{2}}$. Therefore, we can visualize it as in Figure 5.

4. L-Infinity Optimization to Bergman Fans of Matroids

This section begins with the required background on Bergman fans of matroids. Elements of a Bergman fan associated to a matroid $\mathcal{M}$ are called $\mathcal{M}$-ultrametrics as they generalize the ultrametrics of Section 3. The first important result of this section is Proposition 4.5 which says that the set of $\mathcal{M}$-ultrametrics that are l^∞ nearest to a given point forms a tropical polytope. We then define the notion of a "hierarchy of connected flats", which we use to generalize the notion of topology of an ultrametric to arbitrary $\mathcal{M}$-ultrametrics. We then generalize the notions of polytomy and resolution to $\mathcal{M}$-ultrametrics and review Ardila's result generalizing Chepoi and Fichet's algorithm of Theorem 3.1 to $\mathcal{M}$-ultrametrics. These generalizations are used in Algorithm 2 which

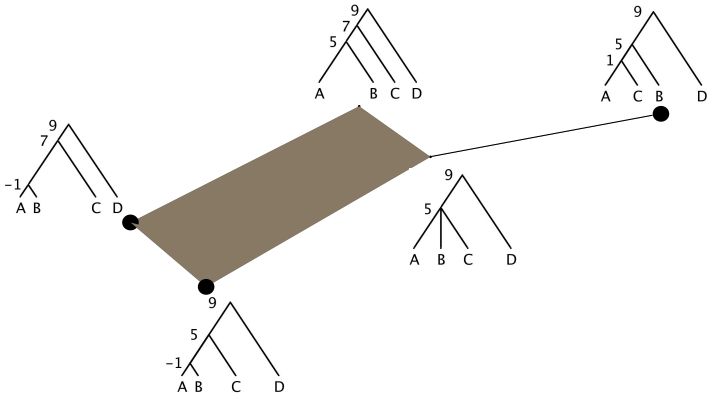

FIGURE 5. The tropical polytope consisting of ultrametrics that are l^∞-nearest to δ. The large points are the tropical vertices.

computes a superset of the tropical vertices of the subset of the Bergman fan containing elements that are l^∞-nearest to x. We also show that two $\mathcal{M}$-ultrametrics have the same topology if and only if they lie in the same cone of the so-called "coarse subdivision" of the Bergman fan associated to $\mathcal{M}$. This section ends with a proposition that suggests an application of $\mathcal{M}$-ultrametrics to phylogenetics when $\mathcal{M}$ is the matroid of a graph.

Familiarity with the fundamental definitions of matroid theory are essential for this section. For this, we refer the reader to [14]. Let $\mathcal{M}$ be a matroid of rank r with ground set E. Each $w \in \mathbb{R}^E$ defines a weight vector on E. Given any basis B of $\mathcal{M}$, the *weight of* B is $\sum_{b\in B} w_b$. If the weight of some basis B is minimal with respect to w, then B is a *w-minimal basis*.

Definition 4.1. Let $\mathcal{M}$ be a matroid on ground set E. If each $e \in E$ appears in some w-minimum basis of $\mathcal{M}$ then w is an *$\mathcal{M}$-ultrametric*. The collection of $\mathcal{M}$-ultrametrics, denoted $\tilde{\mathcal{B}}(\mathcal{M})$, is the *Bergman fan of* $\mathcal{M}$.

Example 4.2. Our recurring example matroid will be $\mathcal{M}(K_4)$, the graphic matroid of the complete graph on four vertices. The ground set E is the collection of edges. Each $x \in \mathbb{R}^E$ is an $\mathbb{R}$-labeling of the edges in $\mathcal{M}(K_4)$. Figure 6 gives two examples of $\mathcal{M}(K_4)$-ultrametrics.

As the name suggests, $\mathcal{M}$-ultrametrics are generalizations of the ultrametrics discussed in Section 3. In particular, the following theorem of Ardila and Klivans tells us that ultrametrics are $\mathcal{M}(K_n)$-ultrametrics where $\mathcal{M}(K_n)$ is the matroid underlying the complete graph on n vertices.

Theorem 4.3 ([3], Theorem 3)**.** *A dissimilarity map on the set $\{1, \ldots, n\}$ is an ultrametric if and only if it is an $\mathcal{M}(K_n)$-ultrametric.*

Ardila identified a strong connection between ultrametric reconstruction using the l^∞-metric and tropical convexity, one component of which is the following.

Proposition 4.4 ([2], Proposition 4.1)**.** *The Bergman fan $\tilde{\mathcal{B}}(\mathcal{M})$ is a tropical polyhedral cone.*

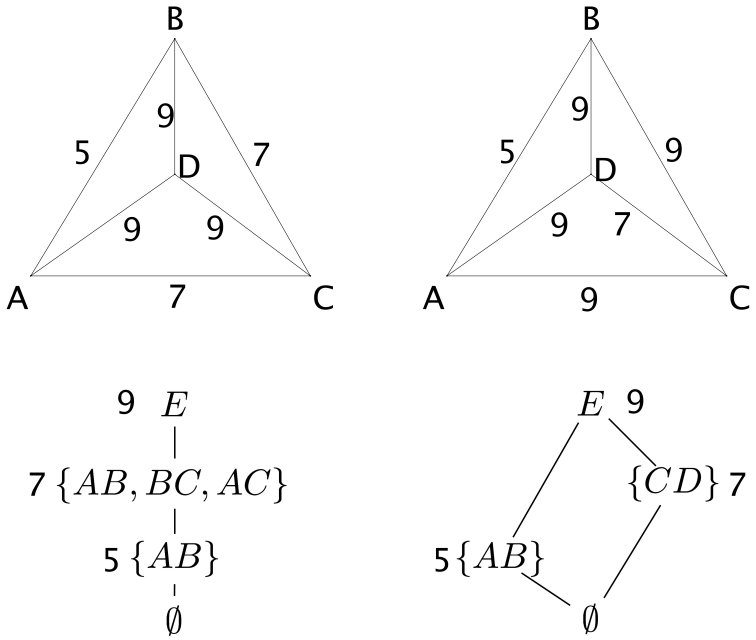

FIGURE 6. $\mathcal{M}(K_4)$-ultrametrics with different topologies

We introduce some notation. Given some $x \in \mathbb{R}^E$, we define the subset of $\tilde{\mathcal{B}}(\mathcal{M})$ consisting of the $\mathcal{M}$-ultrametrics that are l^∞-nearest to x by $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. The next proposition says that this set is a tropical polytope.

Proposition 4.5. *If $\mathcal{M}$ is a matroid on ground set E and $x \in \mathbb{R}^E$, then the subset of the Bergman fan of $\mathcal{M}$ consisting of elements l^∞-nearest to x is a tropical polytope.*

Proof. Let C denote the cube of side-length $d(x, \tilde{\mathcal{B}}(\mathcal{M}))$ centered at x. Therefore we can express

$$C(x, \tilde{\mathcal{B}}(\mathcal{M})) = \tilde{\mathcal{B}}(\mathcal{M}) \cap C.$$

Proposition 4.4 tells us that $\tilde{\mathcal{B}}(\mathcal{M})$ is a tropical polyhedron and C is clearly a tropical polytope. Their intersection is again a tropical polyhedron. Since it is bounded it is by definition a tropical polytope. □

Much of the remainder of this section is devoted to describing the set of tropical vertices of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. We begin by defining a notion of topology for an $\mathcal{M}$-ultrametric which generalizes the rooted tree topology underlying an ultrametric.

We quickly recall the direct sum operation on matroids and matroid connectivity. For more details, see [14, Ch. 4]. Let $\mathcal{M}_1$ and $\mathcal{M}_2$ be matroids on disjoint ground sets E_1 and E_2. We define their direct sum, denoted $\mathcal{M}_1 \oplus \mathcal{M}_2$, to be the matroid on ground set $E_1 \cup E_2$ such that $I \subseteq E_1 \cup E_2$ is independent in $\mathcal{M}_1 \oplus \mathcal{M}_2$ if and only if $I \cap E_i$ is independent in $\mathcal{M}_i$ for each i. A key property is that $F \subseteq E_1 \oplus E_2$ is a flat of $\mathcal{M}_1 \oplus \mathcal{M}_2$ if and only if F_i is a flat of $\mathcal{M}_i$ for each i. It is unfortunate that direct sum and tropical addition are denoted by the same symbol. However, the context will always makes the

proper interpretation clear.

We say that $\mathcal{M}$ is *connected* if $\mathcal{M} = \mathcal{M}_1 \oplus \mathcal{M}_2$ implies that either $\mathcal{M}_1 = \mathcal{M}$ and $\mathcal{M}_2 = \emptyset$ or $\mathcal{M}_2 = \mathcal{M}$ and $\mathcal{M}_1 = \emptyset$. Up to relabeling of indices, each matroid $\mathcal{M}$ can be uniquely written as $\mathcal{M} = \mathcal{M}_1 \oplus \cdots \oplus \mathcal{M}_k$ where each $\mathcal{M}_i$ is connected. Each $\mathcal{M}_i$ is called a *connected component of* $\mathcal{M}$.

Definition 4.6. Let $\mathcal{M}$ be a matroid on ground set E. Let $\mathcal{F}$ be a collection of connected flats of $\mathcal{M}$ partially ordered by inclusion. We call $\mathcal{F}$ a *hierarchy of connected flats* if

(1) $\emptyset \in \mathcal{F}$, and
(2) for all $F, G \in \mathcal{F}$, either $F \subseteq G$, $G \subseteq F$, or $F \cap G = \emptyset$.

Cover relations will be denoted $F \lessdot_{\mathcal{F}} G$, or simply $F \lessdot G$ when $\mathcal{F}$ is clear from context.

As we will see, there is a well-defined way to associate a hierarchy of connected flats to each $\mathcal{M}$-ultrametric. Before proving this, we recall some notation from matroid theory and prove an intermediate lemma. If $\mathcal{M}$ is a matroid on ground set E and $S \subseteq E$, then the *restriction of* $\mathcal{M}$ *to* S, denoted $\mathcal{M}|_S$, is the matroid on ground set S such that $I \subseteq S$ is independent in $\mathcal{M}|_S$ if and only if I is independent in $\mathcal{M}$. The *contraction of* S *in* $\mathcal{M}$, denoted $\mathcal{M}/S$, is the matroid on ground set $E \setminus S$ such that $I \subseteq E \setminus S$ is independent in $\mathcal{M}/S$ if and only if there exists a basis B of $\mathcal{M}|_S$ such that $I \cup B$ is independent in $\mathcal{M}$. When $T \subseteq S \subseteq E$, we will abuse notation and write S/T to mean $\mathcal{M}|_S/T$.

Lemma 4.7. *Let $\mathcal{M}$ be a matroid on ground set E. If $w \in \mathbb{R}^E$ is an $\mathcal{M}$-ultrametric and F is a flat of $\mathcal{M}$, then $w|_F$ is an $\mathcal{M}|_F$-ultrametric.*

Proof. Since w is an $\mathcal{M}$-ultrametric, there is a flag of flats

$$\emptyset = F_0 \subsetneq F_1 \subsetneq \cdots \subsetneq F_k = E$$

such that w has constant value c_i on each $F_i \setminus F_{i-1}$, and $c_1 < \cdots < c_r$. Intersecting all the F_i's with F gives a flag of flats of $\mathcal{M}|_F$ after removing redundant entries. Denote this new flag $\emptyset = G_0 \subsetneq G_1 \subsetneq \cdots \subsetneq G_{k'} = F$. Let $e \in F$. We are going to show that e is contained in some w-minimum basis of $\mathcal{M}|_F$. Let j be the smallest index such that $e \in G_j$. Every $w|_F$-minimum basis of $\mathcal{M}|_F$ can be constructed by choosing a basis of each G_i/G_{i-1} and taking their union. Since G_j is a flat, e is not a loop in G_j/G_{j-1}. Therefore, we can choose a basis for G_j/G_{j-1} containing e and complete it to a $w|_F$-minimum basis of $\mathcal{M}|_F$. □

Proposition 4.8. *Let $\mathcal{M}$ be a matroid on ground set E. Let $w \in \mathbb{R}^E$ be an $\mathcal{M}$-ultrametric. Then there exists a unique hierarchy of connected flats $\mathcal{F}$ such that for each $F \in \mathcal{F}$,*

(1) *w is constant on $F \setminus \left(\bigcup_{G \lessdot F} G\right)$, and*
(2) *if $G \lessdot F$, $g \in G$, and $f \in F$, then $w_g < w_f$.*

Proof. We proceed by induction on the number of connected components of $\mathcal{M}$ and the number of distinct values of w. If w has only one distinct value then the connected components of $\mathcal{M}$ form the unique hierarchy of connected flats satisfying the desired properties.

Now assume $\mathcal{M}$ is connected and that w takes multiple values on E. Let $e \in E$ such

that $w_e = \max_{f \in E} w_f$. Then $F := \{f \in E : w_f < w_e\}$ is a flat of $\mathcal{M}$ since w is an $\mathcal{M}$-ultrametric. So $w|_F$ is an $\mathcal{M}|_F$-ultrametric by Lemma 4.7. By induction on the number of distinct values that w takes, $\mathcal{M}|_F$ has a unique hierarchy of connected flats, $\mathcal{F}$, such that the two desired properties are satisfied on F. Then $\mathcal{F} \cup \{E\}$ is the unique hierarchy of connected flats satisfying the desired properties for $\mathcal{M}$.

Now assume $\mathcal{M}$ is disconnected with $\mathcal{M} = \mathcal{M}_1 \oplus \mathcal{M}_2$ and each $\mathcal{M}_i$ nonempty. Let E_i denote the ground set of $\mathcal{M}_i$. Then E_i is a flat of $\mathcal{M}$ so $w|_{E_i}$ is a $\mathcal{M}_i$-ultrametric by Lemma 4.7. By induction on the number of connected components, $\mathcal{M}_i$ has a unique hierarchy of connected flats $\mathcal{F}_i$ that satisfies the desired properties. Therefore $\mathcal{F}_1 \cup \mathcal{F}_2$ satisfies the desired properties for $\mathcal{M}$ and is also unique. □

Definition 4.9. Let $\mathcal{M}$ be a matroid on ground set E and let $w \in \mathbb{R}^E$ be an $\mathcal{M}$-ultrametric. The unique hierarchy of connected flats from Proposition 4.8 is called the *topology of* w, denoted $\mathcal{T}(w)$.

The terminology above might be unsettling to some readers since it appears to have nothing to do with topology in the usual sense. We use it because it generalizes the notion of tree topology of an ultrametric in the phylogenetics sense (see Proposition 4.31 and paragraph before Definition 4.30). The use of the word "topology" for ultrametrics may be equally unsettling to some, but its use is standard in the phylogenetics literature and so we stick with it.

Given $F \in \mathcal{T}(w)$, we abuse notation and let w_F denote the constant value taken by w_e when $e \in F \setminus \left(\bigcup_{G \lessdot F} G\right)$. We can express an $\mathcal{M}$-ultrametric w by drawing the Hasse diagram for $\mathcal{T}(w)$ and labeling each flat F by w_F.

Example 4.10. Figure 6 shows two $\mathcal{M}(K_4)$-ultrametrics and their topologies. Each flat F is labeled by the constant value taken by elements of $F \setminus \left(\bigcup_{G \lessdot F} G\right)$.

Definition 4.11. Let $\mathcal{M}$ be a matroid on ground set E and let $\mathcal{F}$ be a hierarchy of connected flats of $\mathcal{M}$. If $F \in \mathcal{F}$ such that $\operatorname{rank}(F/\bigcup_{G \lessdot F} G) > 1$, then we say that F is a *polytomy*.

Polytomies for $\mathcal{M}$-ultrametrics generalize the notion of polytomies for ultrametrics. That is, an ultrametric has a polytomy as defined in Section 3 if and only if it has a polytomy as an $\mathcal{M}(K_n)$-ultrametric as defined above. We can resolve a polytomy of an $\mathcal{M}$-ultrametric just as we resolve polytomies in the phylogenetic tree associated to an ultrametric.

Definition 4.12. The *resolutions of* $\mathcal{F}$, denoted $R(\mathcal{F})$, is the collection of all polytomy-free hierarchies of connected flats of $\mathcal{M}$ that contain $\mathcal{F}$ as an induced sub-poset. We let $\mathscr{R}(\mathcal{F})$ denote the collection of flats of $\mathcal{M}$ that appear in some resolution of $\mathcal{F}$ - that is $\mathscr{R}(\mathcal{F}) := \{F : F \in \mathcal{G} \text{ for some } \mathcal{G} \in R(\mathcal{F})\}$.

Example 4.13. The posets in Figure 7 are hierarchies of connected flats of $\mathcal{M}(K_4)$. The hierarchy on the left has a polytomy since $\operatorname{rank}(E/\{AB\}) = 2$. To its right are the three possible resolutions.

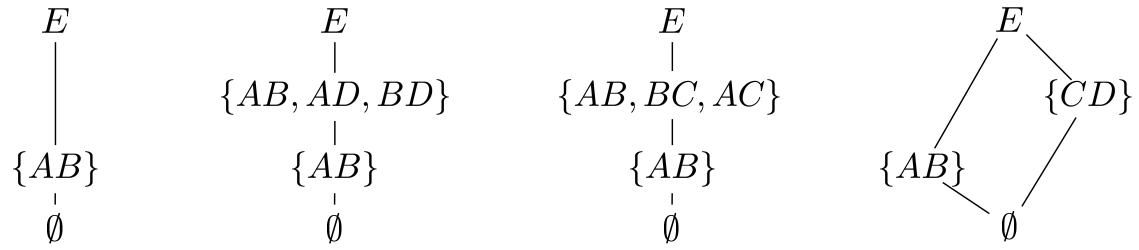

FIGURE 7. A polytomy and its resolutions

This concept of topology is used in our algorithm for computing a superset of the tropical vertices of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. Now we recall the concept of a subdominant $\mathcal{M}$-ultrametric, the existence of which was proven by Ardila in [2].

Definition 4.14. Let $\mathcal{M}$ be a matroid on ground set E and let $x \in \mathbb{R}^E$. Let $x^{\mathcal{M}}$ denote the unique coordinate-wise maximum $\mathcal{M}$-ultrametric which is coordinate-wise at most x. We call $x^{\mathcal{M}}$ the *subdominant $\mathcal{M}$-ultrametric of x*.

Given some $x \in \mathbb{R}^E$, Ardila shows how the first three steps of the algorithm from Theorem 3.1 can be extended to compute a $\mathcal{M}$-ultrametric $x^{\mathcal{M}}$ such that $x^{\mathcal{M}} \leq x$ and is maximal with respect to this property. This ultrametric $x^{\mathcal{M}}$ is called the *subdominant $\mathcal{M}$-ultrametric of x*. Then the subdominant ultrametric can easily be shifted to obtain an l^∞-nearest ultrametric that is coordinate-wise maximal among all l^∞-nearest ultrametrics. This construction can be used to show there exists some $e \in E$ such that u_e is constant across all $u \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. For what follows, $\mathbf{1}$ denotes the all-1s vector in $\mathbb{R}^E$.

Lemma 4.15. *Let $\mathcal{M}$ be a matroid on ground set E, $x \in \mathbb{R}^E$, and $\delta = \frac{1}{2}d(x, x^{\mathcal{M}})$. Define $S := \{e \in E : x_e - x_e^{\mathcal{M}} = 2\delta\}$. Then*

1. *The l^∞-distance from x to $\tilde{\mathcal{B}}(\mathcal{M})$ is δ*
2. *$x^{\mathcal{M}} + \delta \cdot \mathbf{1}$ is an l^∞-nearest $\mathcal{M}$-ultrametric to x*
3. *If y is an l^∞-nearest $\mathcal{M}$-ultrametric, then for all $e \in S$, $y_e = x_e^{\mathcal{M}} + \delta$.*
4. *$x^{\mathcal{M}} + \delta \cdot \mathbf{1}$ is maximal among $\mathcal{M}$-ultrametrics l^∞-nearest to x.*

Proof. The existence of $x^{\mathcal{M}} + \delta \cdot \mathbf{1}$ shows that $d(x, \tilde{\mathcal{B}}(\mathcal{M})) \leq \delta$. Suppose there exists $w \in \tilde{\mathcal{B}}(\mathcal{M})$ such that $d(x, w) < \delta$, then $w - d(x, w) \cdot \mathbf{1}$ is coordinate-wise less than x. Then there exists $e \in E$ such that $x_e - x_e^{\mathcal{M}} = 2\delta$ and so $x_e^{\mathcal{M}} < w_e - d(x, w) \leq x_e$. Thus, $w - d(x, w) \cdot \mathbf{1}$ is an ultrametric coordinate-wise less than x but not coordinate-wise less than $x^{\mathcal{M}}$, contradicting that $x^{\mathcal{M}}$ is the subdominant $\mathcal{M}$-ultrametric. So (1) is proven and (2) immediately follows.

Now we prove (3). Let $y \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. If $y_e < x_e^{\mathcal{M}} + \delta$ for some $e \in S$, then $d(y, x) \geq x_e - y_e = 2\delta - y_e + x_e^{\mathcal{M}} > \delta$ contradicting (1). If $y_e > x_e^{\mathcal{M}} + \delta$ for some $e \in S$, then $y - \delta \cdot \mathbf{1}$ is an $\mathcal{M}$-ultrametric that is coordinate-wise at most x, but $(y - \delta \cdot \mathbf{1})_e > x_e^{\mathcal{M}}$ contradicting that $x^{\mathcal{M}}$ is the subdominant $\mathcal{M}$-ultrametric.

If (4) were false and there existed some $\mathcal{M}$-ultrametric $y \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ such that $y \geq x^{\mathcal{M}} + \delta \cdot \mathbf{1}$ with inequality somewhere, then $y - \delta \cdot \mathbf{1}$ would not be coordinate-wise

less than $x^{\mathcal{M}}$. However, it would be coordinate-wise less than x, thus contradicting that $x^{\mathcal{M}}$ is the subdominant $\mathcal{M}$-ultrametric. □

Definition 4.16. Given $x \in \mathbb{R}^E$, we denote by x^m the l^∞-nearest ultrametric $x^{\mathcal{M}} + d(x, \tilde{\mathcal{B}}(\mathcal{M})) \cdot \mathbf{1}$ and call it the *maximal closest $\mathcal{M}$-ultrametric to x*.

Example 4.17. Letting E denote the edge set of K_4, Figure 8 shows an element of $\mathbb{R}^E$ along with $x^{\mathcal{M}}$ and x^m.

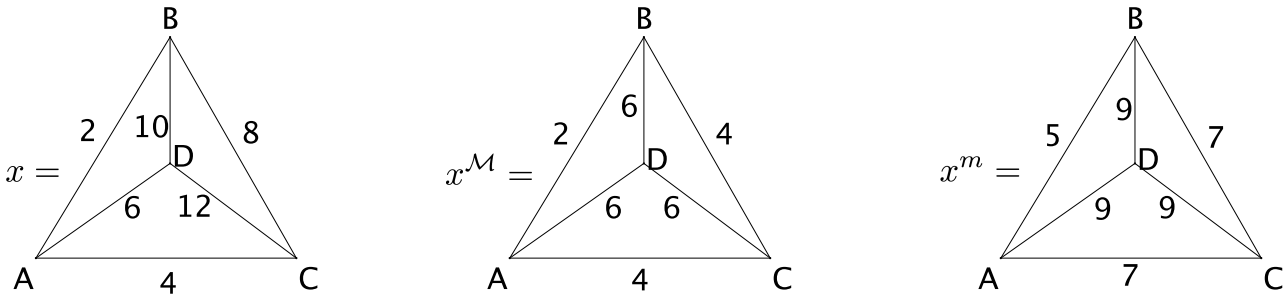

FIGURE 8. A point in $\mathbb{R}^E$ that is not an $\mathcal{M}(K_4)$-ultrametric, its subdominant ultrametric, and its maximal closest ultrametric

Definition 4.18. Let $u \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. We define the *slack of u* to be the following nonnegative vector

$$s(u, x) := u - x + d(x, \tilde{\mathcal{B}}(\mathcal{M})) \cdot \mathbf{1}.$$

In other words, $s(u, x)_e$ gives the maximum amount u_e can be decreased before being further away than $d(x, \tilde{\mathcal{B}}(\mathcal{M}))$ from x in the l^∞-norm.

Definition 4.19. Fix a resolution $\mathcal{F} \in R(\mathcal{T}(u))$ of $\mathcal{T}(u)$. We define the slack of any $F \in \mathcal{F}$ to be

$$s(u, x, \mathcal{F})_F := \min\{s(u, x)_e : e \in F \text{ and } e \notin G \text{ for all } G \lessdot_{\mathcal{F}} F\}.$$

We say that F is *mobile in* $\mathcal{F}$ if $s(u, x, \mathcal{F})_F > 0$ and one of the following holds.

(1) F lies in $\mathcal{T}(u)$ and is not a polytomy, or
(2) $F \notin \mathcal{T}(u)$ and if $G \lessdot_{\mathcal{F}} F$ then $G \in \mathcal{T}(u)$.

Observation 4.20. *Let $F \in \mathscr{R}(\mathcal{T}(u))$. If F is mobile in some resolution containing F then the collection of flats $\{G : G \lessdot_{\mathcal{F}} F\}$ is independent of the choice of resolution. In particular, if F is mobile in some resolution of $\mathcal{T}(u)$ then F is mobile in all resolutions.*

In light Observation 4.20, it makes sense to say that some $F \in \mathscr{R}(\mathcal{T}(u))$ is mobile without reference to any particular resolution. Moreover, when F is mobile, we can write $G \lessdot F$ to mean $G \lessdot_{\mathcal{F}} F$ for some resolution $\mathcal{F} \in R(\mathcal{T}(u))$. Because of this, $s(u, x, \mathcal{F})_F$ does not depend on $\mathcal{F}$ for mobile F, and so in this case we simply write $s(u, x)_F$. When F is not mobile, we say that F is *immobile*.

Example 4.21. Let E be the edge set of K_4 and let x be as in Figure 8. Figure 9 shows the slack vector $s(x^m, x)$ and the topology $\mathcal{T}(x^m)$ of the maximal closest ultrametric where each flat F is labeled by $s(x^m, x)_F$. Note that the flat E is immobile, whereas the other two nonempty flats are mobile.

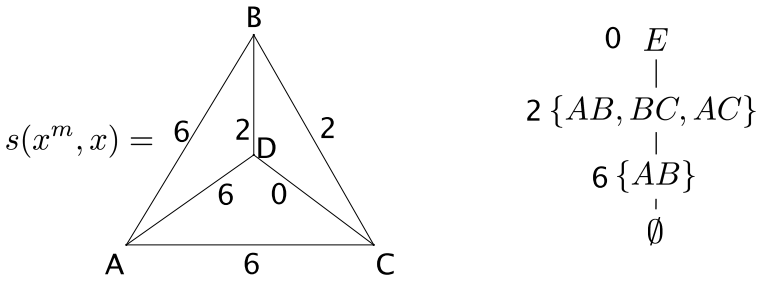

FIGURE 9. Slack vector and topology of x^m

We are now ready to state our algorithm for computing a finite set whose tropical convex hull is $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$, the set of $\mathcal{M}$-ultrametrics that are l^∞-nearest to x. This is given in Algorithm 2.

Algorithm 2 Computes a finite set whose tropical convex hull is $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$

Input: A matroid $\mathcal{M}$ on ground set E and a point $x \in \mathbb{R}^E$
  Compute $x^{\mathcal{M}}$ using Ardila's algorithm [2] and translate it into x^m
  Set $V := \emptyset$, $S_0 := \{x^m\}$, $S_i := \emptyset$ for $i = 1, 2, \ldots$
  for $i := 1, 2, \ldots$ **do**
    if $S_{i-1} = \emptyset$ **then**
      break
    end if
    for $u \in S_{i-1}$ **do**
      for each mobile $F \in \mathscr{R}(\mathcal{T}(u))$ **do**
        initialize $w \in \mathbb{R}^E$ such that $w_e = u_e$ for all $e \notin F \setminus \bigcup_{G \lessdot F} G$
        if $u_F - s(u, x)_F \geq \max\{w_G : G \lessdot F\}$ **then**
          set $w_e = u_F - s(u, x)_F$ for each $e \in F \setminus \bigcup_{G \lessdot F} G$.
        else
          set $w_e = \max\{w_G : G \lessdot F\}$ for each $e \in F \setminus \bigcup_{G \lessdot F} G$
        end if
        put w in S_i
      end for
    end for
  end for
  return $\bigcup_{i=0}^{\infty} S_i$

Before proving correctness of Algorithm 2, we use it to compute an example.

Example 4.22. Let $\mathcal{M} = \mathcal{M}(K_4)$, let x be as in Figure 8, and let $u_1, \ldots, u_5$ be as in Figure 10. Running Algorithm 2 with input $\mathcal{M}$ and x gives $S_0 = \{u_1\}$, $S_1 = \{u_2, u_3\}$, and $S_2 = \{u_4, u_5\}$. Note that u_5 gets added to S_2 two times - once when $u = u_2$ and $F = \{AB\}$ and once when $u = u_3$ and $F = \{AB, AC, BC\}$. Note that this is merely a rephrasing of Example 3.3 in the language of this section.

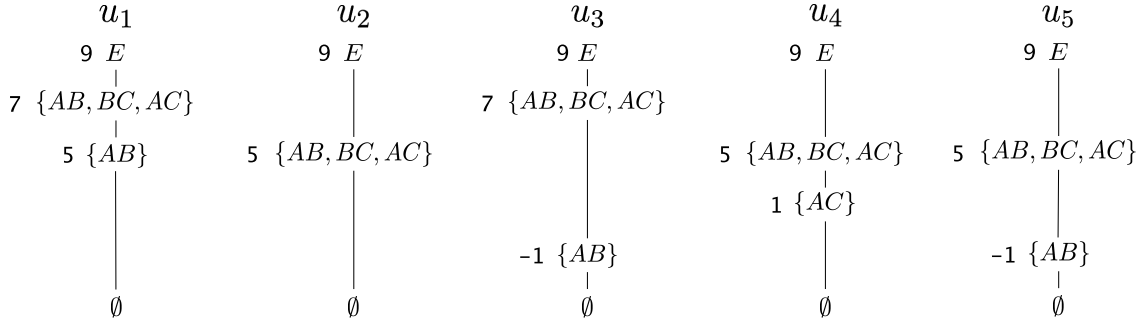

FIGURE 10. $\mathcal{M}(K_4)$-ultrametrics produced by Algorithm 2

Theorem 4.23. *Given a matroid $\mathcal{M}$ on ground set E and $x \in \mathbb{R}^E$, Algorithm 2 returns the tropical vertices of the set of all $\mathcal{M}$-ultrametrics that are l^∞-nearest to x.*

Proof. We first prove termination. Each w produced in the inner for-loop of Algorithm 2 has coordinates that are either $w_e = x^m_f$ or $w_e = x^m_f - s(x^m_f, x)_f$ for some $f \in E$ (note e and f may be different). In particular, there are only finitely many such values. Note that each $w \in S_{i+1}$ is coordinate-wise less than some $w' \in S_i$ with strict inequality at some coordinate. Therefore $S_i = \emptyset$ for some $i > 0$ and so Algorithm 2 terminates.

We now show that each tropical vertex is reached. To that effect let v be a tropical vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. We construct a sequence $x^m = w^0 \geq w^1 \geq \cdots \geq w^k = v$ such that $w^i \in S_i$. Assuming x^i has been constructed, we show how to construct w^{i+1}.

First assume $\mathcal{T}(w^i) \subseteq \mathcal{T}(v)$ with $v \neq w^i$. Since $w^i \geq v$, there exists a flat $F \in \mathscr{R}(\mathcal{T}(v)) \subseteq \mathscr{R}(\mathcal{T}(w^i))$ such that F is mobile in w^i. Moreover we claim that with proper choice of F, we may define $w^{i+1} \in S_{i+1}$ so that $w^{i+1}_e = v_e$ for $e \in F \setminus \bigcup_{G \lessdot_{\mathcal{T}(v)} F} G$ and $w^{i+1} = w^i$ otherwise. As defined, w^{i+1} will be an $\mathcal{M}$-ultrametric if and only if $v_F \geq w^i_G$ for all $G \lessdot_{\mathcal{T}(v)} F$. Choosing such an F is always possible and w^i will be included in S_i in the "if" clause of the innermost for-loop if and only if F is immobile in v. So now assume the only F mobile in w^i satisfying $v_F \geq w^i_G$ for all $G \lessdot_{\mathcal{T}(v)} F$ is also mobile in v. Define $u \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ such that $u_e = v_e - s(v_e, x)$ for all $e \in F \setminus \bigcup_{G \lessdot_{\mathcal{T}(v)} F} G$ and $u_e = v_e$ otherwise. Since F is mobile, $u \neq v$. Then we can write $v = ((v_F - w^i_F) \odot w^i) \oplus u$ implying that v is not a vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$.

Now assume $\mathcal{T}(w^i) \not\subseteq \mathcal{T}(v)$. Define $v^t := (t \odot w^i) \oplus v$. Since $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ is tropically convex, $v^t \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ whenever $t \leq 0$. For each flat $G \in \mathcal{T}(w^i) \setminus \mathcal{T}(v)$, there exists some $t < 0$ such that $G \notin \mathcal{T}(v^t)$. Define t_G to be the maximal such t. Fix G so that t_G is maximized and set $t := t_G$. Define $H \in \mathcal{T}(w^i) \cap \mathcal{T}(v^t)$ minimal such that $G \subsetneq H$. Then $w^i_H > w^i_G \geq v^t_H$. If H is a polytomy in w^i choose $F \subset H$ mobile. Otherwise define

$F := H$. Either way we define w^{i+1} such that $w^{i+1}_e = w^i_G$ whenever $e \in F \setminus \bigcup_{I \lessdot_{\mathcal{T}(w^i)} F} I$ and $w^{i+1}_e = w^i_e$ otherwise. Then w^{i+1} gets included in S_{i+1} in the "else" clause of the innermost for-loop of Algorithm 2. □

Algorithm 2 outputs more than just the tropical vertices of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. However, we will see that if $u \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ is a tropical vertex, then it can have at most one mobile flat. First we need two technical lemmas about tropical polyhedra.

Lemma 4.24 ([1], Proposition 3.1)**.** *Let $P \subseteq \mathbb{R}^E$ be a tropical cone. Then $g \in P$ lies on an extreme ray if and only if there exists some $e \in E$ such that for each $x \in P$, if $x_e = g_e$ and $x \leq g$, then $x = g$.*

Lemma 4.25. *Let $P \subseteq \mathbb{R}^E$ be a tropical polytope with a fixed coordinate - i.e. there exists some $e \in E$ and a constant k such that $x_e = k$ for all $x \in P$. Then* $\mathrm{tcone}(P)$ *is a tropical polyhedral cone. If $v \in \mathrm{tcone}(P)$ with $v_e = k$, then $v \in P$. In this case moreover, v lies on an extreme ray of* $\mathrm{tcone}(P)$ *if and only if v is a vertex of P.*

Proof. Define $C_P := \{(\lambda \odot x, \lambda) : x \in P, \lambda \in \mathbb{R}_{max}\}$. Then C_P is a tropical polyhedral cone, and v is a vertex of P if and only if $(v, 0)$ lies on an extreme ray of C_P (cf. [9, Corollary 2.15]). Let V denote the set of vertices of P. Note that if $(x, x_0) = \bigodot_{v \in V} a_v \odot (v, 0)$ is an arbitrary point of C_P, then $x = \bigodot_{v \in V} a_v \odot v$ lies in $\mathrm{tcone}(P)$. Moreover, we can invert this projection map. Namely, if $x = \bigodot_{v \in V} a_v \odot v$ is an arbitrary point in $\mathrm{tcone}(P)$, then $x_e = k + \max_{v \in V} a_v$ and so $(x, x_e - k) \in C_P$. Therefore if $x \in \mathrm{tcone}(P)$ with $x_e = k$, then $x \in P$. Moreover x is a vertex of P if and only if $(x, 0)$ lies on an extreme ray of C_P which is true if and only if x lies on an extreme ray of $\mathrm{tcone}(P)$. □

Proposition 4.26. *Let v be a tropical vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$. Then at most one flat in $\mathscr{R}(\mathcal{T}(v))$ is mobile.*

Proof. If $F_1, F_2 \in \mathscr{R}(\mathcal{T}(v))$ are distinct mobile flats, then we can define u^i to be the ultrametric obtained by $u^i_e = v_e - \varepsilon$ whenever $e \in F_i \setminus \bigcup_{G \lessdot F} G$ and $u^i_e = v_e$ otherwise. Since F is mobile, $u^i \in C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ for small enough ε. Note that for both $i = 1, 2$ we have $u^i \leq v$. However, for any $e \in E$, $u^i_e = v_e$ will be satisfied for some i. Therefore v cannot lie on a tropical extreme ray of $\mathrm{tcone}(C(x, \tilde{\mathcal{B}}(\mathcal{M})))$ by Lemma 4.24 and therefore it is not a tropical vertex of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ by Lemmas 4.25 and 4.15(3). □

Example 4.27. In Example 4.22, the only outputted ultrametrics that have at most one mobile flat are u_3, u_4, and u_5. Proposition 4.26 therefore implies that the tropical vertices of $C(x, \tilde{\mathcal{B}}(\mathcal{M}(K_4)))$ are contained among these. However, the tropical vertices cannot be a proper subset as the dimension of $C(x, \tilde{\mathcal{B}}(\mathcal{M}(K_4)))$ is two whereas the tropical convex hull of two points has dimension one.

The following proposition gives us a way to use Algorithm 2 to check that all elements of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ have the same topology. Namely, it suffices to check that all $\mathcal{M}$-ultrametrics in the output of Algorithm 2 have the same topology.

Proposition 4.28. *The set of $\mathcal{M}$-ultrametrics that have a particular topology $\mathcal{F}$ is tropically convex.*

Proof. It is clear that topology of a $\mathcal{M}$-ultrametric is preserved under tropical scalar multiplication, so it suffices to show that topology is preserved under tropical sums. To this end, let u, v be $\mathcal{M}$-ultrametrics with topology $\mathcal{F}$. Note that $(u \oplus v)$ is constant on $F \setminus \left(\bigcup_{G \lessdot F} G\right)$ for each $F \in \mathcal{F}$ and that $(u \oplus v)_F > (u \oplus v)_G$ whenever $G \lessdot F$. Proposition 4.8 then implies that $\mathcal{T}(u \oplus v) = \mathcal{F}$. □

Corollary 4.29. *Let $\mathcal{M}$ be a matroid on ground set E and let $x \in \mathbb{R}^E$. Then set of all $\mathcal{M}$-ultrametrics that are l^∞-nearest to x have the same topology if and only if all tropical vertices of $C(x, \tilde{\mathcal{B}}(\mathcal{M}))$ have the same topology.*

We now show that the polyhedral subdivision of $\tilde{\mathcal{B}}(\mathcal{M})$ induced by topology coincides with the coarsest possible subdivision of $\tilde{\mathcal{B}}(\mathcal{M})$. This is useful from a phylogenetics perspective because two $\mathcal{M}(K_n)$-ultrametrics lie in the same open cone in this subdivision if and only if they have the same topology as ultrametrics [3]. Thus our definition of $\mathcal{M}$-ultrametric topology indeed generalizes the concept of ultrametric topology.

Definition 4.30. Let $\mathcal{M}$ be a matroid on ground set E and let $w \in \mathbb{R}^E$. Then $\mathcal{M}_w$ denotes the matroid on ground set E whose bases are the w-minimum bases of $\mathcal{M}$.

The collection of $\mathcal{M}$-ultrametrics that induce a particular collection of w-minimal bases $\mathcal{M}_w$ forms an open polyhedral cone. Their closures form the coarsest possible subdivision of $\tilde{\mathcal{B}}(\mathcal{M})$ (cf. [10, Section 3]). This coincides with the subdivision given by topology.

Proposition 4.31. *For any $w, u \in \tilde{\mathcal{B}}(\mathcal{M})$, $\mathcal{T}(w) = \mathcal{T}(u)$ if and only if $\mathcal{M}_w = \mathcal{M}_u$.*

Proof. We proceed by showing that $\mathcal{T}(u)$ uniquely determines $\mathcal{M}_u$ and vice versa. To see that $\mathcal{T}(u)$ determines $\mathcal{M}_u$, note that each u-minimum basis of $\mathcal{M}$ is obtained by choosing a basis of $F/\left(\bigcup_{G \lessdot F} G\right)$ for each $F \in \mathcal{T}(u)$ and taking their union.

We now show that $\mathcal{M}_u$ determines $\mathcal{T}(u)$. Choose some $\mathcal{M}$-ultrametric w such that $\mathcal{M}_w = \mathcal{M}_u$. Then w determines a flag of flats

$$\mathcal{F} := F_0 \subsetneq F_1 \subsetneq \cdots \subsetneq F_k$$

such that $F_0 = \emptyset$, $F_k = E$, w is constant on each $F_i \setminus F_{i-1}$, and if $e \in F_i \setminus F_{i-1}$ and $f \in F_{i-1}$, then $w_e > w_f$. If each F_i is connected, then $\mathcal{F} = \mathcal{T}(w)$ by Proposition 4.8. Now assume this is not the case and let i be the largest number such that F_i is disconnected. Express F_i as a direct sum $F_i = C \oplus D$. Then we can express the first $i+1$ parts of $\mathcal{F}$ as

$$\emptyset \subsetneq C_1' \oplus D_1' \subsetneq \cdots \subsetneq C_{i+1}' \oplus D_{i+1}'.$$

Without changing $\mathcal{M}_w$, we can change the values of w so that $D_j' = \emptyset$ for all $j \leq \operatorname{rank}(C)$. This is true because any $\mathcal{M}_w$-minimum basis of $\mathcal{M}$ must contain a basis of each C_j'/C_{j-1}' and a basis of each D_j'/D_{j-1}'. Such a change also preserves the ordering of coordinates of w within each connected component, and hence $\mathcal{T}(w)$. Now we can write the first $i+1$ parts of $\mathcal{F}$ as

$$\emptyset \subsetneq C_1 \subseteq \cdots \subsetneq C \subsetneq C \oplus D_1 \subsetneq \cdots \subsetneq C \oplus D.$$

We repeat this process recursively on each disconnected C_j and D_j until every summand is connected. The resulting collection of connected flats is $\mathcal{T}(w)$ by Proposition 4.8. □

When $\mathcal{M} := \mathcal{M}(G)$ is the matroid underlying some graph G, then Algorithm 2 has potential use for phylogenetics even when G is not the complete graph. In particular, it sometimes happens that only a subset of the pairwise distances between n species can be computed within a reasonable budget. Then one may ask the question of which partial ultrametrics are l^∞-nearest to the observed distances. Assuming that the observed distances correspond to the edge set E of a graph G, the following proposition tells us that the above question is equivalent to: given some partial dissimilarity map $x \in \mathbb{R}^E$, which $\mathcal{M}(G)$-ultrametrics are l^∞-nearest to x?

Proposition 4.32. *Let $E \subsetneq \binom{[n]}{2}$, let G be the graph with vertex set $[n]$ and edge set E, and let $x \in \mathbb{R}^E$. Then we may extend x to some ultrametric $x' \in \tilde{\mathcal{B}}(\mathcal{M}(K_n))$ if and only if x is a $\mathcal{M}(G)$-ultrametric.*

Proof. First let x be a $\mathcal{M}(G)$-ultrametric. Let $e \in \binom{[n]}{2} \setminus E$. Let G' be the graph obtained by adding e to G. We can extend x to a $\mathcal{M}(G')$-ultrametric x' by setting x'_e to be the maximum of all the minimum edge weights appearing in some cocircuit of $\mathcal{M}(G')$. That this is indeed an $\mathcal{M}(G')$-ultrametric follows from Ardila's characterization of $\mathcal{M}$-ultrametrics in terms of $\mathcal{M}$'s cocircuits [2]. By induction it follows that x may be completed to an $\mathcal{M}(K_n)$-ultrametric.

Now let $x \in \mathbb{R}^E$ and assume that there exists some $x' \in \tilde{\mathcal{B}}(\mathcal{M}(K_n))$ such that $x_e = x'_e$ for each $e \in E$. Since x' is an $\mathcal{M}(K_n)$-ultrametric, each $e \in E$ appears in some x'-minimal basis of $\mathcal{M}(K_n)$. As $x'_e = x_e$ for each $e \in E$, it follows that each $e \in E$ appears in some x-minimal basis of $\mathcal{M}(G)$. Therefore x is an $\mathcal{M}(G)$-ultrametric. □

5. Example on a biological dataset

In this section we apply our algorithm to the dataset displayed in Figure 11. It consists of pairwise immunological distances between the species dog, bear, raccoon, weasel, seal, sea lion, cat, and monkey that were obtained by Sarich in [15]. It is used in [8] to illustrate the UPGMA and neighbor joining algorithms, which are two other distance-based methods for phylogenetic reconstruction.

	dog	bear	raccoon	weasel	seal	sea lion	cat	monkey
dog	0	32	48	51	50	48	98	148
bear	32	0	26	34	29	33	84	136
raccoon	48	26	0	42	44	44	92	152
weasel	51	34	42	0	44	38	86	142
seal	50	29	44	44	0	24	89	142
sea lion	48	33	44	38	24	0	90	142
cat	98	84	92	86	89	90	0	148
monkey	148	136	152	142	142	142	148	0

Figure 11. Pairwise immunological distances between eight species

Applying Algorithm 1 and using Proposition 4.26 to extract a subset of the output containing the tropical vertices gives us the twenty ultrametrics displayed in Table 1.

Four different tree topologies appear. However, they all have the same tree topology when restricted to monkey, cat, weasel, and dog. It is interesting to note that none of these trees have the same topology as the tree returned by either UPGMA or neighbor joining, and those two trees also have differing topologies (see pages 166 and 168 in [8]).

TABLE 1. A set of ultrametrics whose tropical convex hull is the set of ultrametrics l^∞-nearest to the dataset in Figure 11

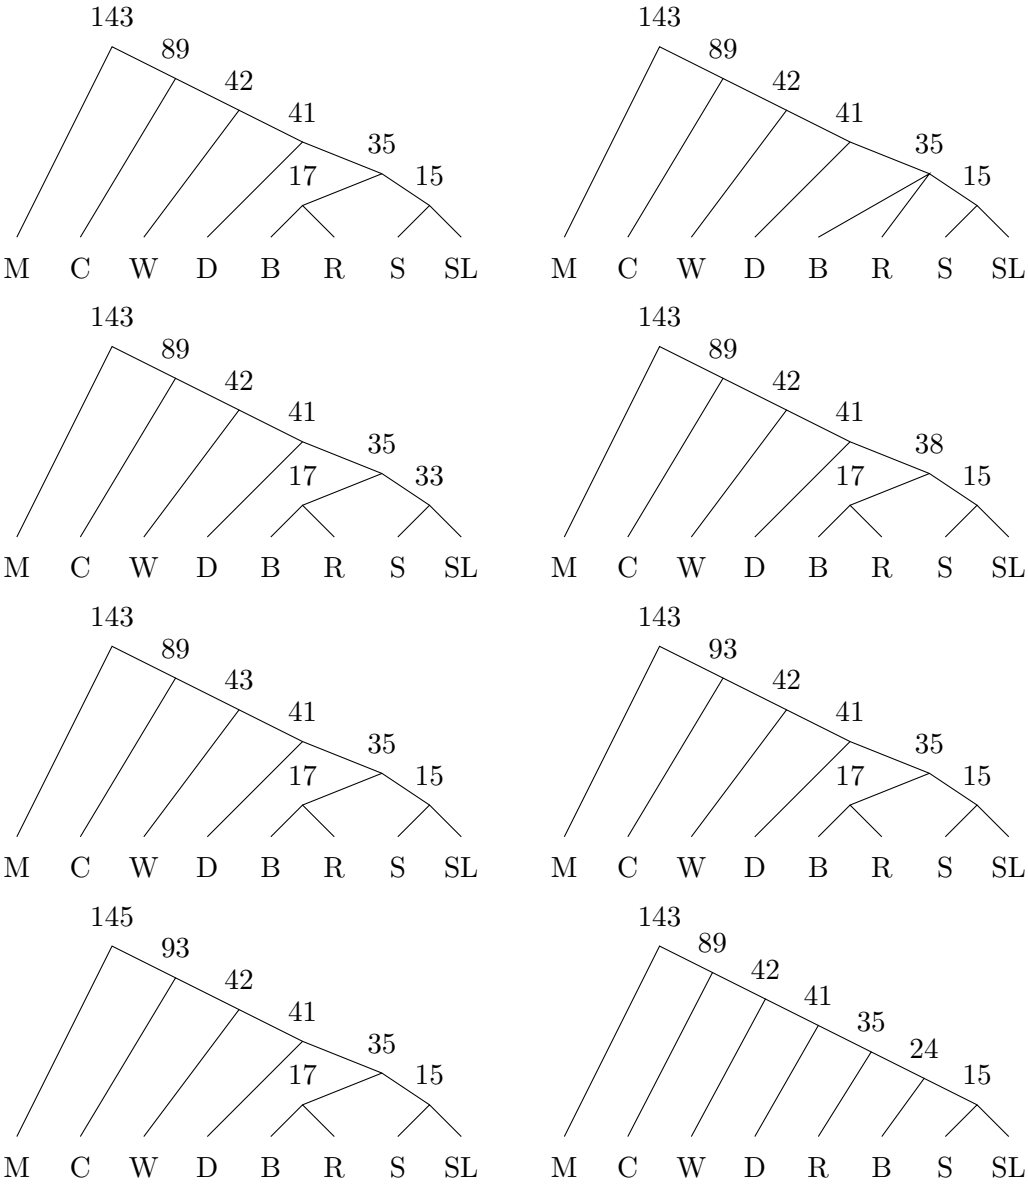

Continued on next page

Table 1 – continued from previous page

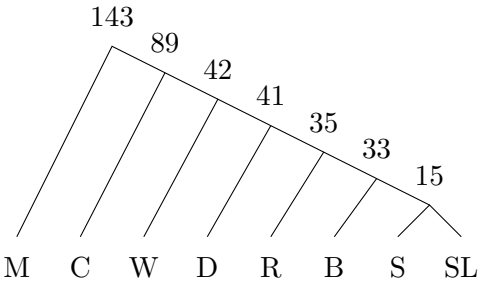

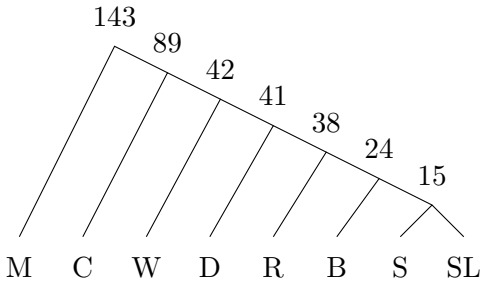

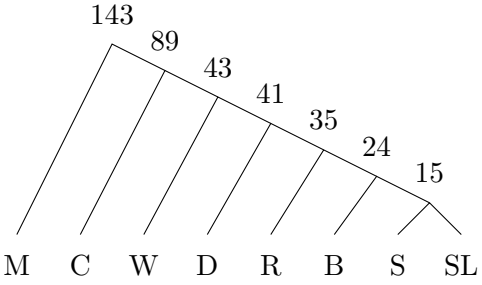

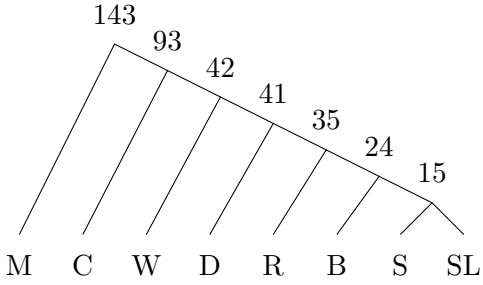

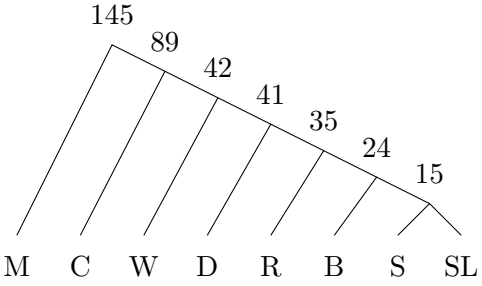

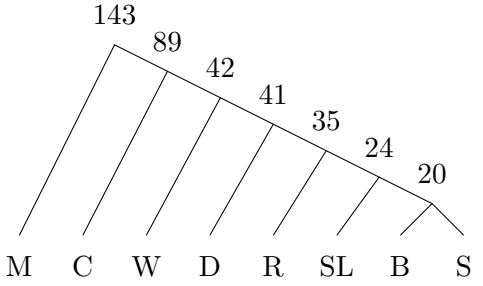

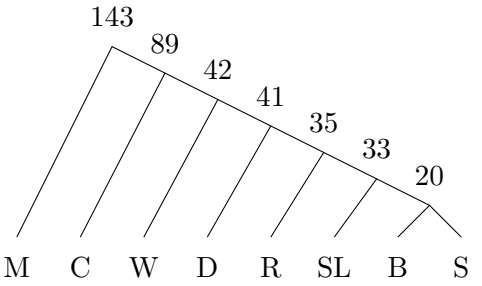

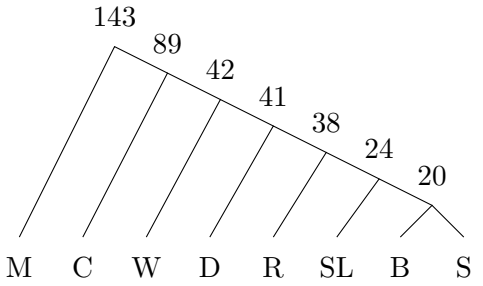

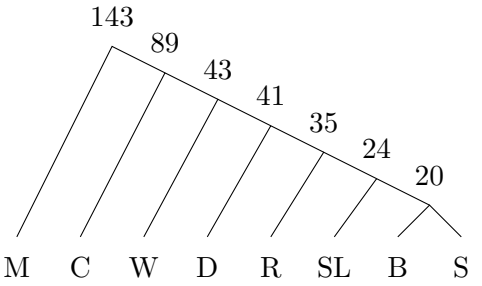

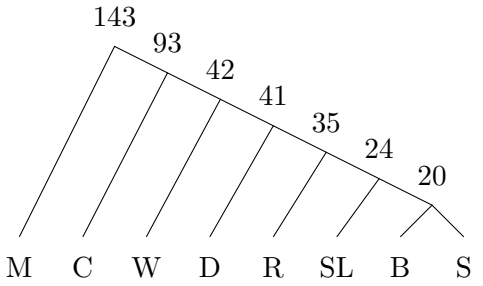

Continued on next page

Table 1 – continued from previous page

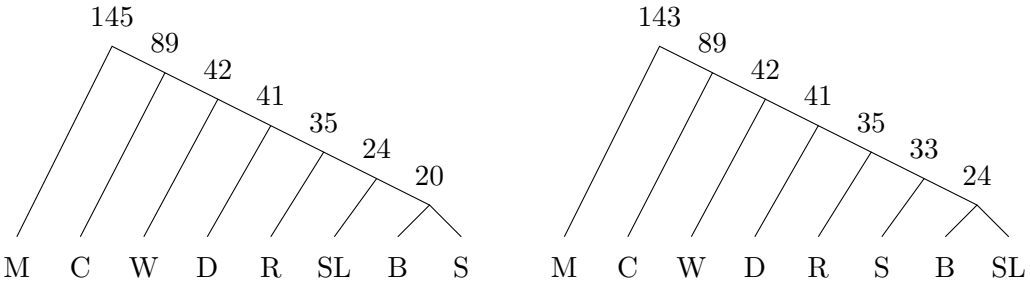

References

[1] Xavier Allamigeon, Stéphane Gaubert, and Éric Goubault. The tropical double description method. In *STACS 2010: 27th International Symposium on Theoretical Aspects of Computer Science*, volume 5 of *LIPIcs. Leibniz Int. Proc. Inform.*, pages 47–58. Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern, 2010.

[2] Federico Ardila. Subdominant matroid ultrametrics. *Annals of Combinatorics*, 8:379–389, 2004.

[3] Federico Ardila and Caroline J. Klivans. The Bergman complex of a matroid and phylogenetic trees. *Journal of Combinatorial Theory, Series B*, 96(1):38 – 49, 2006.

[4] Daniel Irving Bernstein and Colby Long. L-infinity optimization to linear spaces and phylogenetic trees. *arXiv preprint arXiv:1702.05127*, 2017.

[5] V. Chepoi and B. Fichet. l_∞-approximation via subdominants. *Journal of Mathematical Psychology*, 44:600–616, 2000.

[6] Mike Develin and Bernd Sturmfels. Tropical convexity. *Documenta Mathematica*, 9:1–27, 2004.

[7] Martin Farach, Sampath Kannan, and Tandy Warnow. A robust model for finding optimal evolutionary trees. In *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing*, pages 137–145. ACM, 1993.

[8] Joseph Felsenstein and Joseph Felenstein. *Inferring phylogenies*, volume 2. Sinauer associates Sunderland, 2004.

[9] Stéphane Gaubert and Ricardo D Katz. The Minkowski theorem for max-plus convex sets. *Linear Algebra and its Applications*, 421(2):356–369, 2007.

[10] Simon Hampe. *Algorithmic aspects of tropical intersection theory.* PhD thesis, Technische Universität Kaiserslautern, 2014.

[11] Michael Joswig. Tropical halfspaces. *Combinatorial and computational geometry*, 52:409–431, 2005.

[12] Bo Lin, Bernd Sturmfels, Xiaoxian Tang, and Ruriko Yoshida. Convexity in tree spaces. *https://arxiv.org/abs/1510.08797*, 2016.

[13] Bo Lin and Ruriko Yoshida. Tropical Fermat-Weber points. *https://arxiv.org/abs/1604.04674*, 2016.

[14] James Oxley. *Matroid Theory.* Oxford University Press, second edition, 2011.

[15] Vincent M Sarich. Pinniped phylogeny. *Systematic Biology*, 18(4):416–422, 1969.

[16] David Speyer and Bernd Sturmfels. The tropical Grassmannian. *Adv. Geom.*, 4:389–411, 2004.

Department of Mathematics, North Carolina State University, Raleigh, NC 27695
E-mail address: dibernst@ncsu.edu